

Node Sampling using Drifting Random Walks*

Andrés SEVILLA¹, Alberto MOZO², and Antonio FERNÁNDEZ ANTA³

¹ Dpto Informática Aplicada, U. Politécnica de Madrid, Madrid, Spain
 asevilla@eui.upm.es

² Dpto Arquitectura y Tecnología de Computadores, U. Politécnica de Madrid, Madrid, Spain
 amozo@eui.upm.es

³ Institute IMDEA Networks, Madrid, Spain
 antonio.fernandez@imdea.org

Abstract. Sampling a large network with a given probability distribution has been identified as a useful operation. In this paper we propose a distributed algorithm for sampling networks, so that nodes are selected at a special node, called the *source*, with a given probability distribution. This algorithm is based on a new class of random walks, that we call *Drifting Random Walks* (DRW). A DRW starts at the source and *always* moves away from it.

We propose a DRW algorithm for connected networks that selects a node with any desired probability distribution. A drawback of this algorithm is that it needs preprocessing. When the probability distribution is distance based (i.e., the probability of selecting a node is a function of its distance to the source), variants of the DRW algorithm without preprocessing are proposed.

A DRW algorithm has the novel key features that (1) it always finishes in a number of hops bounded by the network diameter, and (2) selects a node with the *exact probability distribution*. Furthermore, unlike previous Markovian (e.g., classical random walks and epidemic) approaches, DRW does not need to stabilize and can efficiently be used as a service to obtain multiple independent samples.

1 Introduction

Sampling a large network with a given distribution has been identified as a useful operation. For instance, sampling nodes with uniform probability is the building block of epidemic information spreading [10,9]. Similarly, sampling with a probability that depends on the distance to a given node [4,14] is useful to construct small world network topologies [11,6,3]. Other applications that can benefit from distance-based node sampling are landmark-less network positioning systems like NetICE9 [13], which does sampling of nodes with special properties to assign synthetic coordinates to nodes. Currently, there is an increasing interest in obtaining a representative (unbiased) sample from the users of online social networks [7]. In this paper we propose a distributed algorithm for sampling networks with a desired probability distribution.

* This research was supported in part by the Comunidad de Madrid grant S2009TIC-1692, Spanish MICINN grant TEC2011-29688-C02-01, and National Natural Science Foundation of China grant 61020106002.

Related Work One technique to implement distributed sampling is to use gossiping between the network nodes. Jelasity et al. [9] present a general framework to implement a uniform sampling service using gossip-based epidemic algorithms. Bertier et al. [3] implement uniform sampling and DHT services using gossiping. As a side result, they sample nodes with a distribution that is close to Kleinberg’s harmonic distribution (one instance of a distance-dependent distribution). Another gossip-based sampling service that gets close to Kleinberg’s harmonic distribution has been proposed by Bonnet et al. [4]. However, when using gossip-based distributed sampling as a service, it is shown in [5] that only partial independence between samples can be guaranteed without re-executing the gossip algorithm.

Another popular distributed technique to sample a network is the use of random walks [15]. Most random-walk based sampling algorithms do uniform sampling [1,7], usually having to deal with the irregularities of the network. Sampling with arbitrary probability distributions can be achieved with random walks by weighting the hop probabilities, for instance using Metropolis-Hastings random walks [12,8].

In [14], it was shown how sampling with an arbitrary probability distribution can be done without communication if a uniform sampling service is available. In that work, like in all the previous approaches, the desired probability distribution is reached when the stationary distribution of a Markov process is reached. The number of iterations (or hops of a random walk) required to reach this situation (the warm-up time) depends on the parameters of the network and the desired distribution, but it is not negligible. For instance, in [15] it is found by simulation that, to achieve no more than 1% error, in a torus of 4096 nodes at least 200 hops of a random walk are required for the uniform distribution, and 500 hops are required for a distribution proportional to the inverse of the distance. In the light of these results, Markovian approaches seem to be inefficient to implement a sampling service, specially if multiple samples are desired.

Contributions In this paper we present an efficient distributed algorithm to implement a sampling service. The basic technique used for sampling is a new class of random walks that we call *Drifting Random Walks* (DRW). A DRW starts at a special node, called the *source*, and *always* moves away from it. The sampling process in the DRW algorithm works essentially as follows. A DRW always starts at the source node. When the DRW reaches a node x , the DRW stops at that node with a *stay probability*. If the DRW stops at node x , then x is the node selected by the sampling. If the DRW does not stop at x , it jumps to a neighbor of x . To do so, the DRW chooses only among neighbors that are at a larger distance from the source than x . (The probability of jumping to each of these neighbors is not necessarily the same.)

We propose a DRW algorithm that samples *any* connected network with *any* probability distribution (given as weights assigned to the nodes). The drawback of this approach is that, before starting the sampling, some preprocessing is required. This preprocessing involves building a spanning tree in the network, and performing a flooding and a convergecast over the tree. Additionally, each node has to maintain state data to be used by the DRW. However, this is compensated by the facts that, once the preprocessing is completed, multiple independent samplings with the exact desired distribution can be efficiently performed, each taking at most D hops (where D is the diameter of the spanning tree).

When the probability distribution is distance-based and nodes are at integral distances (measured in hops) from the source, variants of the DRW algorithm without pre-processing nor state data are proposed. In a *distance-based probability distribution* all the nodes at the same distance of the source node are selected with the same probability. (Observe that the uniform and Kleinberg's harmonic distributions are special cases of distance-based probability distributions.) In these networks, each node at distance $k > 0$ from the source has neighbors (at least) at distance $k - 1$. We can picture nodes at distance k from the source as positioned on a ring at distance k from the source. The center of all the rings is the source, and the radius of each ring is one unit larger than the previous one. Using this graphical image, we refer the networks of this family as *concentric rings networks*.

Observe that *every* connected network can be seen as a concentric rings network. For instance, by finding the breadth-first search (BFS) tree rooted at the source, and using the number of hops in this tree to the source as distance. This topology can also be imposed in real networks. For instance, consider a radio network in which each node has a fixed position assigned (say, with a GPS). Then, fixing a source node, the nodes in the k th concentric rings can be the nodes whose (Euclidean) distance to the source is in the interval $(k - 1, k]$. If the communication radius is reasonably large, the requirements of the concentric rings topology model will be satisfied.

The first variant of DRW algorithm we propose samples with a distance-based distribution in a network with grid topology. In this network, the source node is at position $(0, 0)$ and lattice (Manhattan) distance is used. This grid contains all the nodes that are at a distance no more than the radius R from the source (the grid has hence a diamond shape). The algorithm we derive assigns a stay probability to each node, that only depends on its distance from the source. However, the hop probabilities depend on the position (i, j) of the node and the position of the neighbors to which the DRW can jump to. We formally prove that the desired distance-based sampling probability distribution is achieved. Moreover, since every hop of the DRW in the grid moves one unit of distance away from the source, the sampling is completed after at most R hops.

We have proposed a second variant of the DRW algorithm that samples with distance-based distributions in concentric rings networks *with uniform connectivity*. These are networks in which all the nodes in each ring k have the same number of neighbors in ring $k - 1$ and the same number in ring $k + 1$. Like the grid variant, this variant is also proved to finish with the desired distribution in at most R hops, where R is the number of rings.

Unfortunately, in general, concentric rings networks have no uniform connectivity. This case is faced by creating, on top of the concentric rings network, an overlay network that has uniform connectivity. In the resulting network, the above DRW variant can be used. We propose a distributed algorithm that, if it completes successfully, builds the desired overlay network. We have found via simulations that this algorithm succeeds in building the overlay network in a large number of cases.

In summary, DRW can be used to implement an efficient sampling service because, unlike previous Markovian (e.g., classical random walks and epidemic) approaches, (1) it always finishes in a number of hops bounded by the network diameter, (2) selects a node with the *exact probability distribution*, and (3) does not need warm-up (stabiliza-

tion) to converge to the desired distribution. In the case that preprocessing is needed, this only has to be executed once, independently on the number of samples taken.

The rest of the paper is structured as follows. In Section 2 we introduce concepts and notation that will be used in the rest of the paper. In Section 3 we present the DRW algorithm for a connected network. In Sections 4 and 5 we describe the DRW algorithm on two concentric rings networks: grids and topologies with uniform connectivity. Finally, in Section 6 we present the simulation based study of the algorithm for concentric rings topologies without uniform connectivity.

2 Definitions and Model

Connected Networks In this paper we only consider connected networks. This family includes most of the potentially interesting networks we can find. In every network, we use N to denote the set of nodes and we assume that there is a special node in the network, called the *source* and denoted by s . We assume that each node $x \in N$ has an associated weight $w(x) > 0$. Furthermore, each node *knows* its own weight. The weights are used to obtain the desired probability distribution p , so that the probability of selecting a node x is proportional to $w(x)$. Let us denote $\eta = \sum_{j \in N} w(j)$. Then, the probability of selecting $x \in N$ is $p(x) = w(x)/\eta$. (In the simplest case, $w(x) = p(x), \forall x$ and $\eta = 1$.)

DRW in Connected Networks As mentioned, in order to use DRW to sample connected networks, some preprocessing is done. This involves constructing a spanning tree in the network and performing a weight aggregation process. After the preprocessing, DRW are used for sampling. A DRW starts from the source, jumping to one of its neighbors in the tree. When the DRW reaches a node $x \in N$, it selects x as the sampled vertex with probability $q(x)$, which we call the *stay probability*. If x is not selected, a neighbor y of x in the tree is chosen, using for that a collection of *hop probabilities* $h(x, y)$. The values of $q(x)$ and $h(x, y)$ are computed in the preprocessing and stored at x . The probability of reaching a node $x \in N$ is called the *visit probability*, denoted $v(x)$.

Concentric Rings Networks We also consider a subfamily of the connected networks, which we call *concentric rings networks*. These are networks in which the nodes of N are at integral distances from s . In these networks, no node is at a distance from s larger than a radius R . For each $k \in [1, R]$, we use $\mathbb{R}_k \neq \emptyset$ to denote the set of nodes at distance k from s , and $n_k = |\mathbb{R}_k|$. These networks can be seen as a collection of concentric rings at distances 1 to R from the source, which is the common center of all rings. For that reason, we call the set \mathbb{R}_k the *ring at distance k* . For each $x \in \mathbb{R}_k$, $\gamma_k(x) > 0$ is the number of neighbors of node x at distance $k - 1$ from s (which is only s itself if $k = 1$), and $\delta_k(x)$ is the number of neighbors of node x at distance $k + 1$ from s (which is 0 if $k = R$).

The concentric rings networks considered must satisfy the additional property that the probability distribution is *distance based*. This means that every node $x \in \mathbb{R}_k$ has the same probability p_k to be selected, for all $k \in [1, R]$. This property allows in the subfamilies defined below to avoid the preprocessing required for connected networks.

Grids A first subfamily of concentric rings networks considered is the grid with lattice distances. In this network, the source is at position $(0, 0)$ of the grid, and it contains all the nodes (i, j) so that $i, j \in [-R, R]$ and $|i| + |j| \leq R$. For each $k \in [1, R]$, the set of nodes in ring k is $\mathbb{R}_k = \{(i, j) : |i| + |j| = k\}$. The neighbors of a node (i, j) are the nodes $(i - 1, j)$, $(i + 1, j)$, $(i, j - 1)$, and $(i, j + 1)$ (that belong to the grid).

Uniform Connectivity The second subfamily considered is formed by the concentric rings networks with *uniform connectivity*. These networks satisfy that

$$\forall k, \forall x, y \in \mathbb{R}_k, \delta_k(x) = \delta_k(y) \wedge \gamma_k(x) = \gamma_k(y). \quad (1)$$

In other words, all nodes of ring k have the same number of neighbors δ_k in ring $k + 1$ and the same number of neighbors γ_k in ring $k - 1$.

DRW in Concentric Rings Networks The behavior of a DRW was already described. In the algorithm that we will present in this paper for concentric rings networks we guarantee that, for each k , all the nodes in \mathbb{R}_k have the same visit probability v_k and the same stay probability q_k . A DRW starts from the source, jumping to one of its neighbors (in the first ring). When it reaches a node $x \in \mathbb{R}_k$, it selects x as the sampled vertex with stay probability q_k . If x is not selected, a neighbor $y \in \mathbb{R}_{k+1}$ of x is chosen.

The desired *distance-based probability distribution* is given by the values p_k , $k \in [1, R]$, where it must hold that $\sum_{k=1}^R n_k \times p_k = 1$. The problem to be solved is to define the stay and hop probabilities so that the probability of a node $x \in \mathbb{R}_k$ is p_k .

Observation 1 *If for all $k \in [1, R]$ the visit v_k and stay q_k probabilities are the same for all the nodes in \mathbb{R}_k , the DRW samples with the desired probability iff $p_k = v_k \cdot q_k$.*

3 Sampling in a Connected Network

In this section, we present a DRW algorithm that can be used to sample any connected network. As mentioned, in addition to connectivity, it is required that each node knows its own weight. A node will be selected with probability proportional to its weight.

DRW Algorithm The DRW algorithm for these networks works as follows.

Building a spanning tree The algorithm for first builds a spanning tree of the network. A feature of the algorithm is that, if several nodes want to act as sources for DRW, they can all share the same spanning tree. Hence only one tree for the whole network has to be built. The algorithm used for the tree construction is not important for the correctness of the DRW algorithm. There are several well known algorithms [2] that can be used to build the spanning tree.

Weight aggregation Once the spanning tree is in place, a source that wants to use DRW for sampling has to trigger a process in which nodes compute and store aggregated weights. This is a preliminary process that, like the construction of the tree, has to be executed only once. It involves flooding the whole tree and collecting data back up the source.

Figure 1 describes the behavior of a source node (left side), and of the rest of the tree nodes (right side), respectively, in the weight aggregation process. Before starting

a DRW, a source node has to make sure that each node obtains the accumulated weights of its subtrees. To achieve that, the source node floods the tree sending a *REQUEST* message (Line 3) to its children. The children of a source node are its neighbors in the tree. The rest of the nodes, when they receive a request message, consider as their parent, with respect to the source s , the sender of the message, and consequently their children are the rest of their neighbors. When a copy of the *REQUEST* message reaches a leaf node –a node without children (Line 16), the node returns its weight to its parent in a *RESPONSE* message (Line 18). Otherwise, when the reached node is not a leaf, the *REQUEST* message is forwarded to its children (Line 19). When a node receives a *RESPONSE* message from one of its children containing the accumulated weight of the child branch (Line 20), it stores this value (Line 21). When this node has received the *RESPONSE* messages from all of its children (Line 22), it adds its own weight and the accumulated weights of its children (Line 23), and it sends a *RESPONSE* message containing this value to its parent (Line 24). At the end, this process stops when the source s receives all the *RESPONSE* messages of its children (Lines 6-9). This process is executed only once. After that, many DRW can start from the source (Lines 10-14).

DRW sampling The spanning tree and the precomputed aggregated weights are used by the DRW to perform the samplings (as many as needed). This process is detailed in Figure 1 for the source node (Lines 10-14), and for the rest of the nodes (Lines 25-31). The length of the DRW is bounded by the diameter D of the tree.

Analysis We show now that the algorithm proposed performs sampling with the desired probability distribution.

Theorem 1. *Each node $x \in N$ is visited by the DRW with probability $v(x) = \frac{T(x)}{\eta}$ and selected by the DRW algorithm with probability $p(x) = \frac{w(x)}{\eta}$.*

Proof. We prove the claim by induction on the number of hops from the source s to node x in the spanning tree. The base case is when the node x is at 1 hop of s (i.e., it is a child of s). Then, x is visited with probability $\frac{T(x)}{\eta}$, since x is chosen by s with this probability in the first hop of the DRW (Line 13). If x is visited, then it is selected with probability $q(x) = w(x)/T(x)$ (Line 26). Then, the probability of being selected is

$$\Pr[\text{select } x] = \frac{T(x)}{\eta} \frac{w(x)}{T(x)} = \frac{w(x)}{\eta}.$$

The induction hypothesis assumes the claim true for a node x at distance i from s . We consider a child y of x , which is at distance $i + 1$ from s .

$$\Pr[\text{visit } y] = v(x) (1 - q(x)) \frac{T(y)}{T(x) - w(x)},$$

where $1 - q(x)$ is the probability of not staying at node x , and $\frac{T(y)}{T(x) - w(x)}$ is the probability to choose the child node y in the next hop of the DRW. Then,

$$v(y) = \frac{T(x)}{\eta} \left(1 - \frac{w(x)}{T(x)}\right) \frac{T(y)}{T(x) - w(x)}$$

```

1  task Tree_Source( $s$ )
2   $DRW\_enabled \leftarrow \text{false}$ 
3  send to children REQUEST( $s$ )
4  when RESPONSE( $s, chld, sum$ ) received
5     $T(chld) \leftarrow sum$ 
6    if received RESPONSE from all children
7    then
8       $\eta \leftarrow \sum_{i \in \text{child}(s)} T(i)$ 
9       $DRW\_enabled \leftarrow \text{true}$ 
10 when DRW_START received
11  wait until  $DRW\_enabled$ 
12  choose a node  $x \in \text{child}(s)$ 
13  with probability  $T(x)/\eta$ 
14  send DRW_MSG ( $s$ ) to  $x$ 

15 task Tree_Node( $x, parent$ )
16 when REQUEST( $s$ ) received
17  if  $x$  is a leaf then
18    send RESPONSE( $s, w(x)$ ) to parent
19  else send REQUEST( $s$ ) to children
20 when RESPONSE( $s, chld, sum$ ) received
21   $T(chld) \leftarrow sum$ 
22  if received RESPONSE from all children then
23     $T(x) \leftarrow w(x) + \sum_{i \in \text{child}(x)} T(i)$ 
24    send to parent RESPONSE( $s, x, T(x)$ )
25 when DRW_MSG( $s$ ) received
26  with probability  $q(x) = w(x)/T(x)$  do
27    select node  $x$  and report to source  $s$ 
28  otherwise
29    choose a node  $y \in \text{child}(x)$ 
30    with probability  $\frac{T(y)}{T(x) - w(x)}$ 
31    send DRW_MSG( $s$ ) to  $y$ 

```

Fig. 1. DRW algorithm for connected networks (left: code for source s ; right: code for node x).

$$v(y) = \frac{T(x)}{\eta} \left(\frac{T(x) - w(x)}{T(x)} \right) \frac{T(y)}{T(x) - w(x)} = \frac{T(y)}{\eta}.$$

and

$$\Pr[\text{select } y] = v(y)q(y) = \frac{T(y)}{\eta} \frac{w(y)}{T(y)} = \frac{w(y)}{\eta}.$$

4 Sampling in a Grid

If the algorithm for connected networks is applied to a grid, given its regular structure, the construction of the spanning tree could be done without any communication among nodes, but the weight aggregation process has to be done as before. However, we show in this section that all preprocessing and the state data stored in each node can be avoided if the probability distribution is based on the distance. DRW sampling process was described in Section 2, and we only redefine stay and hop probabilities.

From Observation 1, the key for correctness is to assign stay and hop probabilities that guarantee visit and stay probabilities that are homogenous for all the nodes at the same distance from the source.

Stay probability For $k \in [1, R]$, the stay probability of every node $(i, j) \in \mathbb{R}_k$ is defined as

$$q_k = \frac{n_k \cdot p_k}{\sum_{j=k}^R n_j \cdot p_j} = \frac{n_k \cdot p_k}{1 - \sum_{j=1}^{k-1} n_j \cdot p_j}.$$

As required by Observation 1, all nodes in \mathbb{R}_k have the same q_k . Note that the probability $q_R = 1$, as one may expect.

Hop probability In the grid, the hops of a DRW increase the distance from the source by one unit. We want to guarantee that the visiting probability is the same for each node at the same distance, to use Observation 1. To do so, we need to observe that nodes (i, j) over the axes (i.e., with $i = 0$ or $j = 0$) have to be treated as a special case, because they can only be reached via a single path, while the others nodes can be reached via several paths. To simplify the presentation, and since the grid is symmetric, we give the hop probabilities for one quadrant only (the one in which nodes have both coordinates non-negative). The hop probabilities in the other three quadrants are similar. The first hop of each DRW chooses one of the four links of the source node with the same probability $1/4$. We have three cases when calculating the hop probabilities from a node (i, j) at distance k , $0 < k < R$, to node (i', j') .

- *Case A:* The edge from (i, j) to (i', j') is in one axis. The hop probability of this link is set to $h_k((i, j), (i', j')) = \frac{i+j}{i+j+1} = \frac{k}{k+1}$.
- *Case B:* The edge from (i, j) to (i', j') is not in the axes, $i' = i + 1$, and $j' = j$. The hop probability of this link is set to $h_k((i, j), (i+1, j)) = \frac{2i+1}{2(i+j+1)} = \frac{2i+1}{2(k+1)}$.
- *Case C:* The edge from (i, j) to (i', j') is not in the axes, $i' = i$, and $j' = j + 1$. The hop probability of this link is set to $h_k((i, j), (i, j+1)) = \frac{2j+1}{2(i+j+1)} = \frac{2j+1}{2(k+1)}$.

It is easy to check that the hop probabilities of a node add up to one.

Analysis In the following we prove that the DRW that uses the above stay and hop probabilities selects nodes with the desired sample probability.

Lemma 1. *All nodes at the same distance $k > 0$ to the source have the same visit probability v_k .*

Proof. The proof uses induction. The base case is $k = 1$, which is trivial since the probability of visiting each of the four nodes at distance 1 from the source s is $v_i = 1/4$. Assuming that all nodes at distance $k > 0$ have the same visit probability v_k , we prove the case of distance $k + 1$. Recall that the stay probability is the same q_k for all nodes at distance k .

The probability to visit a node $x = (i', j')$ at distance $k + 1$ depends on whether x is on an axis or not. If it is in one axis it can only be reached from its only neighbor (i, j) at distance k . This happens with probability (case A),

$$\Pr[\text{visit } x] = v_k(1 - q_k) \frac{i + j}{i + j + 1} = v_k(1 - q_k) \frac{k}{k + 1}.$$

If x is not in an axis, it can be reached from two nodes, $(i' - 1, j')$ and $(i', j' - 1)$, at distance k (Cases B and C). Hence, the probability of reaching x is then

$$\Pr[\text{visit } x] = v_k(1 - q_k) \frac{2(i' - 1) + 1}{2(i' + j')} + v_k(1 - q_k) \frac{2(j' - 1) + 1}{2(i' + j')} = v_k(1 - q_k) \frac{k}{k + 1}.$$

Hence, in both cases the visit probability of a node x at distance $k + 1$ is $v_{k+1} = v_k(1 - q_k) \frac{k}{k + 1}$. This proves the induction and the claim.

Theorem 2. *Every node at distance k from the source is selected with probability p_k .*

Proof. If a node is visited at distance k , it is because no node was selected at distance less than k , since a DRW always moves away from the source. Hence, $\Pr[\exists x \in \mathbb{R}_k \text{ visited}] = 1 - \sum_{j=1}^{k-1} n_j p_j$. Since all the n_k nodes in \mathbb{R}_k have the same probability to be visited (from the previous lemma) and the stay probability is $q_k = \frac{n_k p_k}{\sum_{j=k}^R n_j p_j}$, the probability of selecting a particular node x at distance k from the source is

$$\Pr[\text{select } x] = \left(1 - \sum_{j=1}^{k-1} n_j p_j\right) \frac{1}{n_k} \frac{n_k p_k}{\sum_{j=k}^R n_j p_j} = p_k.$$

Where it has been used that $\sum_{j=1}^R n_j p_j = 1$ and that $(1 - \sum_{j=1}^{k-1} n_j p_j) = \sum_{j=k}^R n_j p_j$.

5 Sampling in a Concentric Rings Network with Uniform Connectivity

In this section we derive a variant of DRW algorithm to sample a concentric rings network with uniform connectivity, where all preprocessing is avoided, and only a small (and constant) amount of data is stored in each node. Recall that uniform connectivity means that all nodes of ring k have the same number of neighbors δ_k in ring $k+1$ and the same number of neighbors γ_k in ring $k-1$.

Distributed algorithm The general behavior of the DRW algorithm for these networks was described in Section 2. In order to guarantee that the algorithm is fully distributed, and to reduce the amount of data a node must know a priori, a node at distance k that sends the DRW to a node in ring $k+1$ piggybacks some information. More in detail, when a node in ring k receives the DRW from a node of ring $k-1$, it also receives the probability v_{k-1} of the previous step, and the values p_{k-1} , n_{k-1} , and δ_{k-1} . Then, it calculates the values of n_k , v_k , and q_k . After that, the DRW algorithm uses the stay probability q_k to decide whether to select the node or not. If it decides not to select it, it chooses a neighbor in ring $k+1$ with uniform probability. Then, it sends to this node the probability v_k and the values p_k , n_k , and δ_k , piggybacked in the DRW.

Figure 2 shows the code of the DRW algorithm. The source s sends the DRW with values $v_0 = 1$, $n_0 = 1$, $p_0 = 0$, and δ_0 . Each node in ring k must only know initially the values δ_k , γ_k and p_k . Observe that n_k (number of nodes in ring k) can be locally calculated as $n_k = n_{k-1} \delta_{k-1} / \gamma_k$. The correctness of this computation follows from the uniform connectivity assumption (Eq. 1).

Analysis The uniform connectivity property can be used to prove by induction that all nodes in the same ring k have the same probability v_k to be reached. The stay probability q_k is defined as $q_k = p_k / v_k$. Then, from Observation 1, the probability of selecting a node x of ring k is $p_k = v_k q_k$. What is left to prove is that the value v_k computed in Figure 2 is in fact the visit probability of a node in ring k .

Lemma 2. *The values v_k computed in Figure 2 are the correct visit probabilities.*

```

1 task  $DRW(x, k, \delta_k, \gamma_k, p_k)$ 
2 when  $(v_{k-1}, p_{k-1}, n_{k-1}, \delta_{k-1})$  received:
3    $n_k \leftarrow n_{k-1} \frac{\delta_{k-1}}{\gamma_k}; \quad v_k \leftarrow n_{k-1} \frac{v_{k-1} - p_{k-1}}{n_k}; \quad q_k \leftarrow \frac{p_k}{v_k}$ 
4   with probability  $q_k$  do select node  $x$  and report to  $s$ 
5   otherwise
6     choose a neighbor  $y$  in ring  $k + 1$  with uniform probability
7     send  $(v_k, p_k, n_k, \delta_k)$  to  $y$ 

```

Fig. 2. Drifting Random Walk algorithm for node x in ring k .

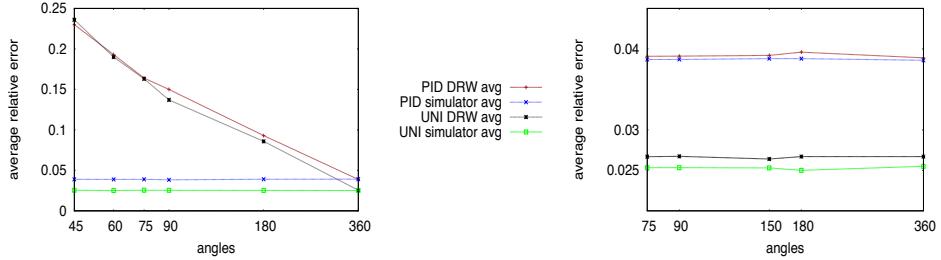


Fig. 3. UNI and PID scenarios without uniform connectivity. Without using the AAP algorithm (left side) and using it (right side).

Proof. Let us use induction. For $k = 1$ the visit probability of a node in ring 1 is $1/n_1$, while the value computed by the algorithm is $v_1 = n_0(v_0 - p_0)/n_1 = 1/n_1$. For a general k , assume the value v_{k-1} is the correct visit probability of ring $k - 1$. The visit probability of a node in ring k is $v_{k-1}n_{k-1}(1 - q_{k-1})/n_k$, which replacing $q_{k-1} = p_{k-1}/v_{k-1}$ yields the expression used in Figure 2 to compute v_k .

The above lemma, together with the previous reasoning, proves the following.

Theorem 3. *Every node at distance k of the source is selected with probability p_k .*

6 Concentric Rings Networks without Uniform Connectivity

Finally, we are interested in evaluating, by means of simulations, the performance of the DRW algorithm variant for concentric rings with uniform connectivity when it is used on a more realist topology: a concentric rings network *without uniform connectivity*. The experiment has been done in a concentric rings topology of 100 rings with 100 nodes per ring, and it places the nodes of each ring uniformly *at random* on each ring. This deployment does not guarantee uniform connectivity. In order to establish the connectivity of nodes, we do a geometric deployment. A node x in ring k is assigned a position in the ring. This position can be given by an angle α . Then, each network studied will have associated a connectivity angle β , the same for all nodes. This means that x will be connected to all the nodes in rings $k - 1$ and $k + 1$ whose position (angle) is in the interval $[\alpha - \beta/2, \alpha + \beta/2]$. We compare the relative error of the DRW

```

1 function AssignAttachmentPoints(x, k)
2   ap  $\leftarrow \frac{LCM(n_k, n_{k+1})}{n_k}$ 
3   C  $\leftarrow \mathbb{N}_{k+1}(x)$  /* neighbors of x in ring k + 1 */
4   Ax  $\leftarrow \emptyset$  /* Ax is a multiset */
5   loop
6     choose c from C
7     send ATTACH_MSG to c
8     receive RESPONSE_MSG from c
9     if RESPONSE_MSG = OK then
10      ap  $\leftarrow ap - 1$ 
11      add c to Ax /* c can be in Ax several times */
12    else C  $\leftarrow C \setminus \{c\}
13  until (ap = 0)  $\vee (C = \emptyset)$ 
14  if (ap = 0) then return Ax
15  else return FAILURE$ 
```

Angle	% success
15°	0%
30°	0%
45°	3%
60°	82%
75°	99%
90°	100%
150°	100%
180°	100%
360°	100%

Fig. 4. Assignment Attachment Points (AAP) Function (left side). Success rate of the AAP algorithm as a function of the connectivity angle (right side).

algorithm when sampling with two distributions: the uniform distribution (UNI) and a distribution proportional to the inverse of the distance (PID). We define the relative error e_i for a node x in a collection C of s samples as $e_i = \frac{|f_{sim_x} - f_x|}{f_x}$, where f_{sim_x} is the number of instances of x in collection C obtained by the simulator, and $f_x = p_x \cdot s$ is the expected number of instances of x with the ideal probability distribution (UNI or PID). We compare the error of the DRW algorithm with the error of a generator of pseudorandom numbers. For each configuration, a collection of 10^7 samples has been done.

Figure 3 (left side) presents the results obtained in the UNI and PID scenarios. In both cases, we can see that the DRW algorithm performs much worse than the UNI and PID simulators. The simulation results show a biased behavior of DRW algorithm because the condition of Eq. 1 is not fulfilled in this experiment (i.e. a node has no neighbors, or there are two nodes in a ring k that have different number of neighbors in rings $k - 1$ or $k + 1$).

AAP Algorithm To eliminate the errors observed when there is no uniform connectivity, we propose a simple algorithm to transform the concentric rings network without uniform connectivity into an overlay network with uniform connectivity.

To preserve the property that the visit probability is the same for all the nodes in a ring, nodes will use different probabilities for different neighbors. Instead of explicitly computing the probability for each neighbor, we will use the following process. Consider rings k and $k + 1$. Let $r = LCM(n_k, n_{k+1})$, where LCM is the *least common multiple* function. We assign $\frac{r}{n_k}$ attachment points to each node in ring k , and $\frac{r}{n_{k+1}}$ attachment points to each node in ring $k + 1$. Now, the problem is to connect each attachment point in ring k to a different attachment point in ring $k + 1$ (not necessarily in different nodes). If this can be done, we can use the algorithm of Figure 2, but when a DRW is sent to the next ring, an attachment point (instead of a neighbor) is chosen uniformly. Since the number of attachment points is the same in all nodes of ring k

and in all nodes of ring $k + 1$, the impact in the visit probability is that it is again the same for all nodes of a ring.

The connection between attachment points can be done with the simple algorithm presented in Figure 4, in which a node x in ring k contacts its neighbors to request available attachment points. If a neighbor that is contacted has some free attachment point, it replies with a response message *RESPONSE_MSG* with value *OK*, accepting the connection. Otherwise it replies to x notifying that all its attachment points have been connected. The node x continues trying until its $\frac{r}{n_k}$ attachment points have been connected or none of its neighbors has available attachment points. If this latter situation arises, then the process failed. Combining these results with the analysis of Section 5, we can conclude with the following theorem.

Theorem 4. *Using attachment points instead of links and the distributed DRW-based algorithm of Figure 2, it is possible to sample a concentric rings network without uniform connectivity with any desired distance-based probability distribution p_k , provided that the algorithm of Figure 4 completes (is successful) in all the nodes.*

Figure 3 (right side) shows the results when using the AAP algorithm. As we can see, the differences have disappeared. The conclusion is that, when nodes are placed uniformly at random and AAP is used to attach neighbors to each node, DRW performs as good as perfect UNI or PID simulators.

In general, the algorithm of Figure 4 may not complete. It is shown in the table of Figure 4 (right side) the success rate of the algorithm for different connectivity angles. It can be observed that the success rate is large as long as the connectivity angles are not very small (at least 60°). (For an angle of 60° the expected number of neighbors in the next ring for each node is less than 17.) For small angles, like 15° and 30° , the AAP algorithm is never successful. For these cases, the algorithm for connected network presented in Section 3 can be used.

References

1. A. Awan, R. A. Ferreira, S. Jagannathan, and A. Grama. Distributed uniform sampling in unstructured peer-to-peer networks. In *HICSS*. IEEE Computer Society, 2006.
2. Cüneyt F. Bazlamaççi and Khalil S. Hindi. Minimum-weight spanning tree algorithms a survey and empirical study. *Comput. Oper. Res.*, 28:767–785, July 2001.
3. M. Bertier, F. Bonnet, A.-M. Kermarrec, V. Leroy, S. Peri, and M. Raynal. D2ht: The best of both worlds, integrating rps and dht. In *EDCC*, pages 135–144. IEEE CS, 2010.
4. F. Bonnet, A.-M. Kermarrec, and M. Raynal. Small-world networks: From theoretical bounds to practical systems. In E. Tovar, Ph. Tsigas, and H. Fouchal, editors, *Proc. OPODIS*, volume 4878 of *Lecture Notes in Computer Science*, pages 372–385. Springer, 2007.
5. Y. Busnel, R. Beraldí, and R. Baldoni. On the uniformity of peer sampling based on view shuffling. *Journal of Parallel and Distributed Computing*, 71(8):1165 – 1176, 2011.
6. P. Fraigniaud and G. Giakkoupis. On the searchability of small-world networks with arbitrary underlying structure. In Leonard J. Schulman, editor, *STOC*, pages 389–398. ACM, 2010.
7. M. Gjoka, M. Kurant, C. T. Butts, and A. Markopoulou. Walking in facebook: A case study of unbiased sampling of osns. In *INFOCOM*, pages 2498–2506. IEEE, 2010.
8. W.K. Hastings. Monte carlo samping methods using markov chains and their applications. *Biometrika*, pages 97–109, 1970.

9. M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. Gossip-based peer sampling. *ACM Trans. Comput. Syst.*, 25(3), 2007.
10. D. Kempe, J. M. Kleinberg, and A. J. Demers. Spatial gossip and resource location protocols. *J. ACM*, 51(6):943–967, 2004.
11. J. M. Kleinberg. Navigation in a small world. *Nature*, 406(6798), August 2000.
12. N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21:1087–1092, June 1953.
13. D. Milić and and T. Braun. Netice9: A stable landmark-less network positioning system. In *Local Computer Networks (LCN), 2010 IEEE 35th Conference on*, pages 96–103, oct. 2010.
14. A. Sevilla, A. Mozo, M. Lorenzo, J. López-Presa, P. Manzano, and A. Fernández Anta. Biased selection for building small-world networks. In *OPODIS 2010*, vol. 6490 of *LNCS*.
15. M. Zhong and K. Shen. Random walk based node sampling in self-organizing networks. *SIGOPS Oper. Syst. Rev.*, 40:49–55, July 2006.